

# A k-NN-based Localization Approach for Crowdsourced Air Traffic Communication Networks

Martin Strohmeier\*, Vincent Lenders<sup>+</sup>, Ivan Martinovic\*

\*University of Oxford, United Kingdom <sup>+</sup>armasuisse, Switzerland

{martin.strohmeier}, {ivan.martinovic}@cs.ox.ac.uk, vincent.lenders@armasuisse.ch

**Abstract**—In this work, we argue that current state-of-the-art methods of aircraft localization such as multilateration are insufficient, in particular for modern crowdsourced air traffic networks with random, unplanned deployment geometry. We propose an alternative, a grid-based localization approach using the k-Nearest Neighbor (k-NN) algorithm, to deal with the identified shortcomings. Our proposal does not require any changes to the existing air traffic protocols and transmitters, and is easily implemented using only low-cost, commercial-off-the-shelf hardware.

Using an algebraic multilateration algorithm for comparison, we evaluate our approach using real-world flight data collected with our collaborative sensor network OpenSky. We quantify its effectiveness in terms of aircraft location accuracy, surveillance coverage, and the verification of false position data.

Our results show that the grid-based k-NN approach can increase the effective air traffic surveillance coverage compared to multilateration by a factor of up to 2.5. As it does not suffer from dilution of precision to the same extent, it is more robust in noisy environments and performs better in pre-existing, unplanned receiver deployments. We further find that the mean aircraft location accuracy can be increased by up to 41% in comparison with multilateration while also being able to pinpoint the origin of potential spoofing attacks conducted from the ground.

**Index Terms**—air traffic control, wireless localization, crowdsourced networks, k-nearest neighbors

## I. INTRODUCTION

Air traffic control (ATC) is the backbone of what is arguably the key means of personal transport in the modern world. As traffic continues to grow dramatically, ATC has to manage ever more aircraft. Large European airports, such as London Heathrow or Frankfurt/Main, experience spikes of more than 1,500 daily take-offs and landings, and industry forecasts predict that world-wide flight movements will double between 2015 and 2034 [1]. Additionally, as Unmanned Aerial Vehicles (UAV) enter the civil airspace, they must learn to co-exist with manned aircraft and existing air traffic control systems. While forecasts project a steady 5% annual increase of global manned flight traffic over the next 20 years, UAV are projected to outgrow traditional air traffic by several orders of magnitude: In 2035, 250,000 UAV are expected to be operating in the US alone, compared to a mere 45,000 passenger aircraft around the globe [2]. However, as this paradigm shift progresses, many technological issues have yet to be solved to ensure the safe control of both manned and unmanned aircraft.

The key issue in controlling the airspace is to know where an aircraft is at any given time. For the pilot, this is achieved

through navigational aids, such as Distance Measuring Equipment (DME) or, more recently, satellite navigation. For ground controllers, the traditional options to obtain an aircraft’s position comprise voice communication via VHF/HF and primary and secondary surveillance radar (PSR and SSR, respectively). Recently, technological developments and stricter separation needs have given rise to other methods of aircraft localization, most notably the Automatic Dependent Surveillance–Broadcast protocol (ADS-B) and multilateration (MLAT).

While ADS-B relies purely on the aircraft correctly broadcasting their own position to other aircraft and ground stations, MLAT can provide an independent means of localization by exploiting the time differences of arrival (TDoA) of signals received at several different ground stations. Both methods achieve greatly increased localization accuracy compared to previous radar surveillance and are key technologies in next generation ATC concepts such as SESAR or NextGen [3].

At the same time, crowdsourced air traffic communication networks have gained importance over the past decade. Large private companies such as FlightRadar24 and PlaneFinder,<sup>1</sup> research networks such as OpenSky,<sup>2</sup> enthusiast websites,<sup>3</sup> and increasingly flight authorities themselves [4] rely less and less on planned deployments of ATC receivers. Instead, they use distributed networks which are randomly deployed. Contrary to traditional, carefully planned receiver networks, this crowdsourced use of cheap sensors provides a number of new challenges to existing localization algorithms.

In this work, we investigate alternatives to multilateration for independent localization in air traffic communication networks. We present a new method based on expected TDoAs, which is robust in noisy environments and depend less on the system’s receiver geometry. We show that it is effective even in unplanned deployments of cheap, crowd-sourced, off-the-shelf ATC receiver networks. Our approach does not require changes to existing technology standards or to the aircraft’s legacy hardware equipment, which is particularly important given aviation’s long adoption and certification cycles [5].

We make the following contributions in this paper:

- We develop a new grid-based method to localize aircraft based on their wireless signals. We utilize a combination of the k-Nearest Neighbor algorithm and the expected

<sup>1</sup><http://flightradar24.com>, <http://planefinder.net>

<sup>2</sup><http://opensky-network.org>

<sup>3</sup><http://adsbexchange.com>

TDoA of ATC signals to estimate their origin.

- We analyse the disadvantages of multilateration, the current state of the art solution for aircraft localization. Concretely, we examine the robustness in noisy environments, the coverage in a typical air traffic surveillance area, and the efficiency of the signal usage.
- We evaluate our approach on real-world data and show that it performs better in wide area settings than multilateration. Compared to the latter, our approach is cheaper and more scalable, and improves surveillance range, detection speed and location accuracy in real-world environments. In particular, our approach is much less dependent on perfect system geometry.
- Finally, we study the effectiveness of verifying air traffic control data. We show that verification of legitimate and false flight data using our approach is quicker and can also provide improved localization of the origin of false data injections on or near the ground.

We have used the core localization idea in previous work for security and data verification purposes [6]. The present work focuses further on localization and provides extensive analysis concerning the benefits in modern crowdsourced air traffic communication sensor networks. We have provided new data and new findings to substantiate our approach, including the theoretical underpinning of expected TDoAs, a robustness analysis, a new discussion of its usefulness for verification, and comparisons with multilateration in crowdsourced networks.

The remainder of this paper is organized as follows. Section II discusses the background necessary for the aircraft localization problem, including the related work. Section III analyzes the problems of multilateration while Section IV illustrates the differences between TDoA-based algorithms. Section V explains the concept of expected TDoAs, and how it can be exploited for lightweight aircraft localization. Section VI details our experimental setup, whereas Section VII evaluates the scheme against real-world flight data. Section VIII examines the scheme’s data verification effectiveness. Finally, Section IX summarizes and concludes this work.

## II. BACKGROUND

### A. Modern Air Traffic Control

Since World War II, ATC surveillance data has been provided by PSR and SSR in various forms. PSR is a term comprising non-cooperative radar localization systems that use a rotating antenna radiating a pulse position-modulated and highly directional electromagnetic beam on a low GHz band. The pulses are returned by all reflecting targets in the airspace; measuring the bearing and round trip time provides the targets’ positions.

In contrast, SSR is a cooperative technology comprising the so-called transponder Modes A, C, and S. Ground stations interrogate the transponders of aircraft in range using digital messages on the 1030 MHz frequency, which reply with the

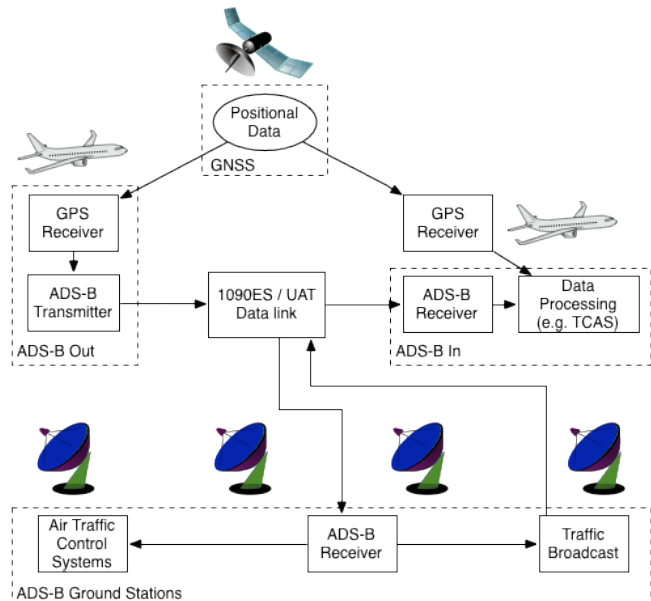


Figure 1: Overview of the ADS-B system architecture. Aircraft receive positional data that is transmitted via the ADS-B Out subsystem over the 1090ES or the UAT data link. It is then received and processed by ground stations and by other aircraft via the ADS-B In subsystem.

requested information on the 1090 MHz channel. SSR provides additional target information compared to PSR besides the position, including identity and altitude.

One of the most recently introduced ATC technologies is the ADS-B protocol, which is the satellite-based successor of SSR. It is currently being rolled out in most airspaces and promises lower cost and more accurate surveillance [7]. Contrary to primary and secondary radar, it does not rely on interrogations. Using two potential datalinks (1090 Extended Squitter, or Universal Access Transceiver), the aircraft automatically broadcasts its position and velocity (twice per second) as well as its identification (every five seconds) to all other aircraft and ground stations in the vicinity (see Fig. 1 for an illustration).

As it is possible to capture all popular wireless protocols with cheap off-the-shelf wireless receivers, crowdsourced sensor networks have increased in popularity over recent years [5]. Using any wireless receiver such as popular software-defined radios, any signal sent by the aircraft, including ADS-B and SSR, can be used for TDoA-based localization such as multilateration [8].

### B. Characteristics of Aircraft Localization

To be able to develop an appropriate localization method, in particular for crowdsourced networks, it is crucial to first identify the characteristics of the ATC environment. The following characteristics distinguish the aircraft localization problem from other wireless localization problems (e.g., wireless sensor networks or vehicular ad hoc networks):

- **Outdoor line-of-sight environment:** Contrary to many localization problems found in academic research, the aircraft location problem is naturally outdoors. On the

1090 MHz channel, the line of sight (LoS) is a crucial factor in receiving signals. We require an outdoor LoS propagation model for our work in terms of loss and propagation.

- **Vast distances:** In *wide area* ATC surveillance, the distances covered are naturally much larger than in more local or indoor problems. Aircraft flying at cruising altitudes (typically 35,000 feet or higher for commercial aircraft) can be observed up to the radio horizon of 400 km or more. This is orders of magnitude larger than typical indoor localization problems.
- **Few multipath effects:** At typical aircraft cruising altitudes, we experience comparably few diffractions leading to multipath effects that influence signal characteristics. This enables us to use simpler theoretical models than in more complex indoor and multipath-rich environments. Most importantly, the propagation timings between aircraft positions and sensors can be approximated easily by using the speed of light  $c$ .

### C. Related Work

Indoor and outdoor localization problems have been studied extensively in the literature, often in the scope of sensor networks and radar applications. Liu et al. [9] give an overview of the techniques used in wireless indoor positioning including the different algorithms (k-Nearest Neighbor, lateration, least squares and Bayesian among others) and primitives such as received signal strength (RSS), TDoA, time of arrival (ToA) and angle of arrival (AoA). RSS-based methods are the most popular within any type of wireless networks as they are often readily supported out of the box and do not require additional hardware such as high precision clocks or antenna arrays.

While TDoA systems are limited in indoor environments (due to multipath effects and non-availability of time synchronization and clocks fine-grained enough to provide good results at very short distances [10]), they offer very superior performance in long-distance outdoor environments such as those encountered in the aircraft localization problem where the non-line-of-sight error is not the dominant error source [11]. Errors in RSS-based outdoor localization are typically much larger, ranging up to a few hundred meters already in comparatively small outdoor settings of 2 km [12].

In terms of algorithms, the k-Nearest Neighbors (k-NN) algorithm has proven to do very well in short-distance, indoor RSS fingerprinting compared to other methods [13], although it has been studied less in long-distance scenarios (such as aircraft localization) and can become computationally more expensive with very large databases.

Overall, the main (distributed) localization approach used within aviation is MLAT [6], which we discuss in detail in the next section and which provides the baseline for our evaluations. To the best of our knowledge the combination of the k-NN algorithm with TDoA as a primitive has not been studied. In the following, we argue that this combination is beneficial in particular for crowdsourced networks with random, imperfect system geometry in which the performance of MLAT suffers strongly.

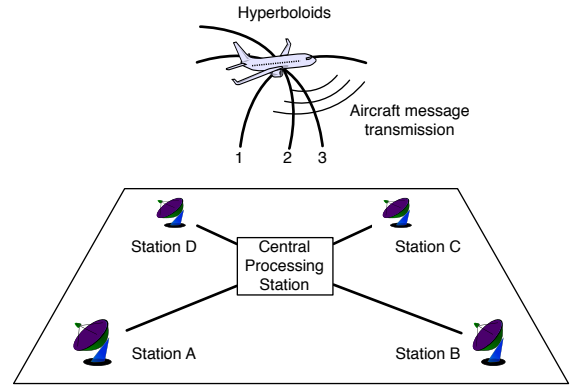


Figure 2: Basic multilateration architecture with four receivers and three resulting hyperboloids.

## III. MULTILATERATION

MLAT is a proven and well-understood concept that is used in civil and military navigation and already serves as a backup for ATC around some airports. It has been the consensus solution in academia and aviation circles as a backup solution for primary and secondary radar [3].

### A. Concept

To localize the 2D-origin of a signal using MLAT, three (or more) receiver stations measure the time at which they receive the same message from an aircraft. They send this data to the central processing station which can calculate the aircraft's position from the intersecting hyperboloids that result from the time difference of arrival between the receiver stations.

Traditionally, the algorithms to solve this problem are classified into open form (or iterative) and closed form (or direct). This classification of algorithms is based on their need of information from external sources: those that require a previous estimation of the solution are considered open form and those that do not require such information are closed form.

Mantilla-Gaviria et al. [14] discuss the different MLAT models used for airport surface surveillance and classify them into three main categories: numerical, statistical and algebraic. The resulting problem obtained by these models is then numerically solved, e.g., by using least squares or singular value decomposition-based regularization.

### B. Downsides

However, despite its popularity, MLAT suffers from a number of known pitfalls.

First, MLAT is highly susceptible to noise, outliers, and even minor measurement errors outside a small core area. An important quality metric for a deployment and its MLAT accuracy with respect to the target object's relative position is the *geometric dilution of precision*, or GDOP (see Fig. 3). It describes the effect of a deployment on the relationship between the errors of the obtained TDoA measurements and their resulting impact on the final errors in the object's calculated position, or formally:

$$\Delta Location Estimate = \Delta Measurements \cdot GDOP \quad (1)$$

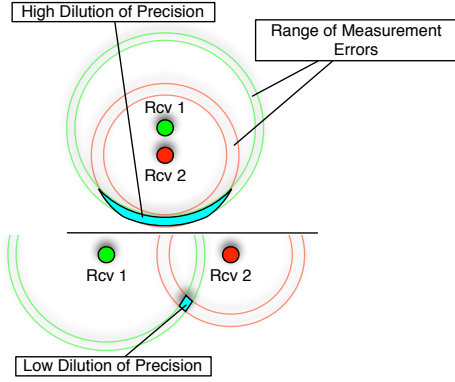


Figure 3: Geometric dilution of precision. The circles show the measurement errors of the respective receivers; the intersections demonstrate the area where the true location of the measured object can be found. An adverse placement of receivers in relation to the target (top) can severely affect the outcome of the localization compared to a favorable deployment (bottom).

GDOP is widely used in positioning systems such as GPS, where good ratings for this multiplier are commonly considered to be below 6, with 10 to be fair and everything over 20 to be of poor quality [15]. For a full explanation of how to calculate different DOP values, see [16].

Second, traditional MLAT systems are very expensive: Where ADS-B needs only a single receiver for accurate wide area surveillance of up to 400 km, MLAT requires every signal to be received by at least three stations with little noise. On top of this, geographical obstacles (e.g., mountain ranges, oceans) make it relatively more difficult to install a comprehensive wide area system at the desired service level. However, even where this is not a problem, MLAT’s discussed dilution of precision characteristics necessitate a well-planned system geometry to accurately cover a given area [14]; an unplanned, or crowdsourced sensor network deployment cannot typically provide good enough GDOP values to achieve the same service quality.

Considering these drawbacks and the fact that modern crowdsourced networks with cheap off-the shelf sensors are becoming more popular, there is a need for other TDoA-based approaches that can improve on these problems.

#### IV. ANALYSING SCALABILITY AND COVERAGE OF CROWDSOURCED LOCALIZATION

One of our main goals is to tackle MLAT’s scalability and coverage problems, especially in deployments with less than perfect system geometry. An ATC data communications network consists of a given number of sensors that are deployed outside, in a line of sight with the airspace they are expected to cover. Naturally, overlapping reception ranges between receivers are required to obtain TDoAs. If more sensors are to receive the same message, they need to be located closer together. While this increases the overlap, it also




Figure 4: The map shows the geometric dilution of precision of our 8 sensor measurement system for wide area multilateration. All sensors are on the ground, the assumed measurement altitude is 38,000 feet.

	Absolute	Relative	Area covered
All messages	63,410,017	100%	100%
# seen by $\geq 2$ sensors	26,305,043	41.48%	46.51%
# seen by $\geq 3$ sensors	8,912,305	14.06%	17.25%
# seen by $\geq 4$ sensors	2,560,727	4.04%	5.28%
# seen by $\geq 5$ sensors	394,362	0.62%	0.84%
# seen by $\geq 6$ sensors	18,777	0.0003%	0.0004%
# seen by $\geq 7$ sensors	97	$1.5 * 10^{-6}\%$	$2.1 * 10^{-6}\%$
# MLAT & GDOP < 10	3,319,618	5.24%	6.67%

Table I: Statistics on the utilized OpenSky dataset. The table shows the absolute and relative number of messages collected by a given amount of sensors. The last column provides the relative area covered by that number of sensors.

decreases the overall ADS-B coverage of the receivers. Worse even, only a small part of the MLAT coverage is usable, since GDOP causes its accuracy to deteriorate quickly. Methods not suffering as much from GDOP and working with fewer sensors could vastly improve efficiency compared to MLAT.

Figure 4 illustrates the dilution of precision experienced by multilateration in our crowdsourced deployment using 8 ADS-B sensors. Covering an area of 3 degrees longitude  $\times$  3 degrees latitude at an altitude of 38,000 feet, it is indicative of a wide area multilateration system used for en-route airspace surveillance. As can be observed, the darkest areas, which illustrate the lowest GDOP values, only make up a fraction of the surveillance area, while the system’s geometry causes most parts to have unusably high dilution of precision.

To demonstrate this fact with real flight data, we analysed more than 60 million ADS-B messages from aircraft at cruising altitudes (ca. 38,000 ft) using 8 OpenSky sensors from






Figure 5: An example illustrating the calculation of Expected TDOAs. The assumed distance of the sender to both receivers is divided by  $c$ . Subtracting the smallest time  $t_i$  from the other times gives the TDOAs relative to receiver  $i$ .

an unplanned and crowdsourced deployment in Switzerland. Table I illustrates the number of messages that are picked up by a given number of receivers. Only 14% of all received messages are seen by 3 or more sensors on the ground and can be used for MLAT. If we only take into account messages with GDOP below 10, we are left with only 5.24% of usable messages.

If we look at the area covered, we find that in our deployment the MLAT coverage makes up roughly 17% of the overall covered area. The area where MLAT is reliably accurate is even smaller, only 6.67% of the total area is covered by messages with a GDOP of less than 10. Thus, the usability of the received ADS-B messages for MLAT is reduced by a factor of more than 2.5 compared to localization methods that would not be affected by GDOP.

## V. DESIGNING A NEW AIRCRAFT LOCALIZATION APPROACH FOR CROWDSOURCED DEPLOYMENTS

We propose a grid-based k-NN approach which is more robust against dilution of precision and works accurately regardless of the system geometry. As laid out above, this characteristic increases the signal utilisation and overall coverage of a deployment, thus vastly reducing costs and improving the detection range and speed of new aircraft at the same time.

### A. Expected TDOAs

The key insight of our approach is the use of *Expected TDOAs*. As outlined before, the time differences of arrival of a received signal between multiple sensors are a physical primitive that can be used to establish the possible location(s) of a sender. For our approach, we extend this concept and (pre-)calculate the time differences that we would *expect* to see based on any given location.

Concretely, we can use an idealized outdoor LoS propagation model suitable for the aircraft localization problem. We then calculate the absolute propagation times of an ATC signal's origin to the receiving ground stations by dividing the distances  $d_1, \dots, d_n$  between the sender and each of the stations  $i = 1, \dots, n$  by the speed of light  $c$  (see Fig. 5).<sup>4</sup>

$$t_i = \frac{d_i}{c} \quad (2)$$

<sup>4</sup>As the propagation is not happening in a vacuum, this is an approximation, however, the difference is insignificant [17].

---

**Algorithm 1** Localization offline phase. Requires coordinates of sensors and grid as input and outputs the training sets for the online phase.

---

```

1: Input: gridcoords, sensors, squaresize
2:
3: trainingset  $\leftarrow$  [ ]
4: grid  $\leftarrow$  constructGrid(gridcoords, squaresize)
5: for  $\forall$  sensorcombinations do
6:   TDoA_training  $\leftarrow$  [ ]
7:   for  $\forall$  square  $\in$  grid do
8:     TDoAs  $\leftarrow$  computeTDoAs(sensors.coords, square)
9:     TDoAtraining.add(TDoAs, gridsquare)
10:  end for
11:  trainingset.add(TDoAtraining,
12:    sensorcombination)
13: end for

```

---



---

**Algorithm 2** Localization online phase. Requires the number of neighbors  $k$  and the *trainingsets* from the offline phase as input and calculates the location estimate of the message's origin.

---

```

1: Input: k, trainingset, flight
2:
3: loop
4:   m  $\leftarrow$  newPositionMessage(flight)
5:   r  $\leftarrow$  receivers(m)
6:   if numberOfReceivers(m) > 2 then
7:     TDoAs  $\leftarrow$  calculateTDoAs(m)
8:     trainingset  $\leftarrow$  getTrainingset(r)
9:     knn  $\leftarrow$  runKNN(trainingset, TDoAs, k)
10:    estimate  $\leftarrow$  getCenter(knn)
11:  end if
12: end loop

```

---

Subtracting the smallest resulting time  $t_{min}$  from the other times provides the Expected TDOAs relative to the nearest receiver  $i$ :

$$ETDoA_i = t_i - t_{min} \quad (3)$$

We use this calculation of Expected TDOAs to estimate the origin of the signal, as explained in the following section.

### B. Aircraft Localization with Expected TDOAs

Putting these findings together, we design a novel approach to locate aircraft. In an offline training phase, a 2D grid of the surveillance area is computed that contains the expected TDoA measurements for each position for the given sensor deployment. In the online phase, for every incoming message, the  $k$  nearest neighbors of the messages' TDOAs on the grid are looked up. These neighbors are then averaged, resulting in the final estimate of the sender's location.

a) *Offline phase*: Over an exemplary grid of  $N \times M$  squares, we calculate the fingerprint vector of expected TDOAs between the deployed sensors for every square using equations 2 and 3. Based on this, we create the final training set

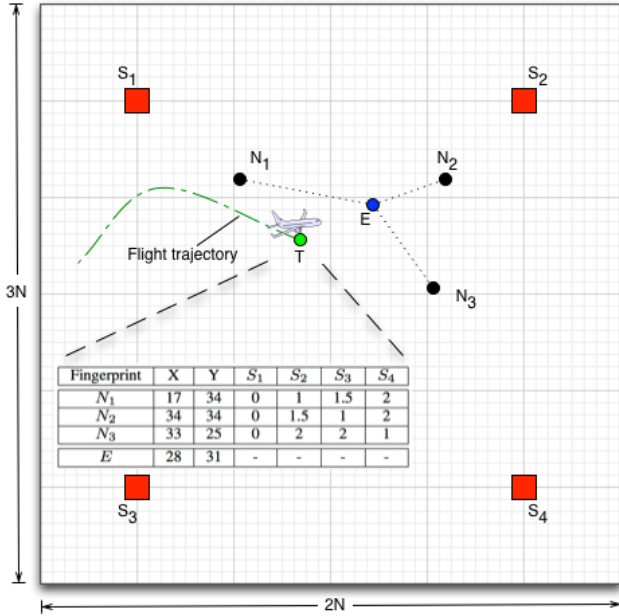


Figure 6: Localization with 3-NN. Using TDoA data from 4 sensors  $S_1, \dots, S_4$ , the 3 nearest neighbors  $N_1, N_2, N_3$  in the lookup table are averaged to obtain the final estimate  $E$ .

by generating every subset of combinations with at least 2 sensors ( $\sum_{i=2}^n \binom{n}{i}$ , with  $n$  being the number of sensors), e.g., 26 sets overall for a 5 sensor deployment. This is required when a new message is received by fewer than the maximum number of deployed sensors, with 2 sensors (or one TDoA measurement) as the lower limit. In that case the appropriate set is chosen during the online phase to find the  $k$  nearest neighbors. Algorithm 1 details the offline training phase.

*b) Online phase:* In the online phase, new message data is analyzed and the location determined (see Algorithm 2 for an overview of the whole process). Using the  $k$ -Nearest Neighbors algorithm, we obtain the closest points from our training grid that match the fingerprints of our test data.

Setting the number of nearest neighbors to  $k$ , we match the received physical fingerprint  $R = TDoA_1, \dots, TDoA_n$  to the saved grid fingerprint  $F$  based on their Euclidean distance

$$D_{(R,F)} = \sqrt{\sum_{i=1}^n (R_{TDoA_i} - F_{TDoA_i})^2} \quad (4)$$

It is intuitive that in the spatial domain of our grid there are multiple neighbors that are approximately the same distance from our point of interest, hence  $k$  is an important parameter influencing the accuracy. If  $k > 1$ , the positions of all  $k$  neighbors are averaged by taking the mean of the longitude and the latitude (see Fig. 6 for an illustration). This constitutes the final estimate of the aircraft position, which on average is closer to the true location than any single neighbor as our evaluation will show.

## VI. EXPERIMENTAL SETUP

We will now describe our experimental setup and design, i.e., how we collected the data used to evaluate our approach,

and the design of the grid used for the localization algorithm. We implement all algorithms in MATLAB.

### A. Data Collection and Hardware

As ADS-B has been in the roll-out phase for years, we can use real-world data to estimate the propagation characteristics of ADS-B messages. We do not make any assumptions on hardware features such as sending power or antennas as there are many configurations found in different aircraft.

For our evaluation, we rely on real-world ADS-B data which we obtained from OpenSky a participatory sensor network for air traffic communication data, specifically ADS-B. It provides access to 3 years of historical raw message data as well as metadata, and offers a very fast query infrastructure, ideal for large-scale research projects. As of October 2016, it has saved more than 200 billion air traffic communication messages, covers about  $3,000,000 km^2$  on three continents, captures more than 16,000 unique aircraft every day, and has seen over 100,000 different aircraft overall. For more detail on OpenSky, its use cases for aviation security and privacy, and its big data infrastructure, refer to [8], [18], [19].

For the present analysis, we use a dataset that spans the period between 26 June 2013 and 25 June 2014. This dataset contains more than 60 million ADS-B messages received from SBS-3 sensors manufactured by Kinetic Avionics. Besides the message content, they provide a timestamp of the message reception. From this data, we use 5 sensors that are closely located together to be able to calculate their TDoA data. The timestamps have a clock resolution of 50 ns. The sensors have omnidirectional antennas and can receive signals from a distance of up to roughly 400 km.

### B. Synchronization

As our low-cost SBS-3 sensors do not provide built-in synchronization (e.g., via GPS), we synchronize our data a posteriori with the help of positional ADS-B messages sent by aircraft. By using the positional information in those messages and approximating their respective propagation time, we can recover the timing offset between our ground station sensors and achieve global synchronization. We also take into account the drift of the internal clocks to improve the results. Overall, this approach enables us to achieve synchronization that is low-cost and works well with minimal requirements. More accurate and efficient synchronization using GPS could help to further improve on the accuracy of our results.

### C. Grid Design

We construct a 2D grid over a typical flight altitude of  $38,000 ft$  (ca.  $11,582 m$ ) with a size of 2 degrees longitude and 2 degrees latitude which, due to the Earth's spherical geometry, translates to an area of ca.  $150 km \times 220 km = 33,000 km^2$ . We obtain evenly-spaced approximate squares where the number of squares (or the squares' size) is a trade-off between performance and accuracy as elaborated in the evaluation section. Of course, computation time and accuracy also depend on the size of the surveillance area.  $33,000 km^2$  are representative for wide area ATC surveillance, covering aircraft's en-route flight phase at cruising altitude.

Horizontal Error [m]	MLAT	600 m <sup>2</sup> Grid	300 m <sup>2</sup> Grid	150 m <sup>2</sup> Grid	75 m <sup>2</sup> Grid	50 m <sup>2</sup> Grid
Mean	199.46	171.01	134.37	122.31	118.14	116.454
Median	91.87	140.38	98.60	84.92	80.38	78.63
RMSE	334.47	225.51	198.14	190.29	187.31	185.79
99th percentile	1306.70	902.08	870.18	870.61	841.33	835.63
Relative comp. time	62.3%	100%	399%	1599%	7272%	16375%

Table II: Horizontal errors in different grid square sizes using k-NN vs. MLAT, with 5 sensors and  $k = 5$ . k-NN shows a better mean accuracy than MLAT of up to 41% in our dataset.

## VII. EVALUATION

In this section, we use the collected flight data to verify our approach. We analyze the location accuracy by comparing it to the GPS data broadcast by aircraft using ADS-B and discuss the impact of the number of neighbors and the grid setup on its performance. We further compare it against an algebraic MLAT algorithm as described in [20],<sup>5</sup> to show its improved efficiency in the examined aircraft localization setting.

### A. Aircraft Localization Accuracy

To ensure a baseline for the accuracy of the localization method, we compare it with the GPS-based ADS-B position claims of legitimate flight data. We use a part of the whole dataset (comprising over 100,000 positional ADS-B messages), where every message has been seen by 5 sensors, providing us with sufficient TDoA measurements for our comparative analysis. All location claims are within the pre-defined surveillance grid in terms of latitude and longitude, while their mean altitude is 11,148.8 m ( $\sigma = 687.59$  m).

Table II shows the localization quality using k-NN with squares of five different sizes over an area of 33,000 km<sup>2</sup> with  $k = 5$ . As expected, increasing the number of squares has a positive impact; the smaller the square, the more accurate location predictions become. For example, a reduction in grid square size from 600 m<sup>2</sup> to 300 m<sup>2</sup> improves mean accuracy by 37.5%. This naturally comes with a trade-off as the computational time to run the k-NN algorithm increases linearly by 400%. Overall, we found that 150 m<sup>2</sup> provides a good trade-off between accuracy and performance.

Concerning the optimal choice of  $k$ , Fig. 7 illustrates the accuracy gained by averaging a higher number of neighbors. We can see a large improvement until  $k = 5$  especially for a grid square size of 600 m<sup>2</sup>. Further decreases in mean accuracy are small and much less pronounced with smaller square sizes.

We compare k-NN with an algebraic MLAT algorithm using the same TDoA measurements from 5 sensors. The results show that with a 600 m<sup>2</sup> grid size, k-NN does 14.2% better than MLAT on mean errors, increasing to 41% for a 50 m<sup>2</sup> grid size. Overall, we find that k-NN does better than MLAT on noisy TDoA measurements such as those we experienced in our real-world data. Especially the more outlier-sensitive metrics RMSE and mean improve with k-NN while MLAT generally shows good median results. Since k-NN is more robust against dilution of precision, this is to be expected as the mean GDOP in our dataset is 24.35 ( $\sigma = 8.06$ ). Taking

<sup>5</sup>Algebraic approaches require more receivers, but provide better accuracy and lower computational resources than numerical models [14].

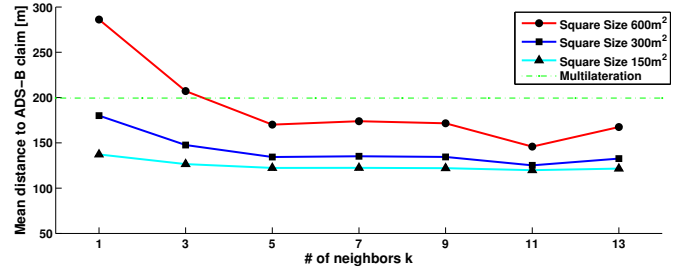


Figure 7: Optimal choice of neighbors  $k$  for different square sizes (MLAT as comparison).

Error [m]	MLAT	2 Sens.	3 Sens.	4 Sens.	5 Sens.
Mean	199.5	26,956.7	311.8	147.3	122.3
Median	91.9	22,737.1	145.4	95.8	84.9
RMSE	334.5	33,380.4	761.3	237.6	190.3
99th tile	1306.7	63,500.2	2,469.6	983.7	870.6

Table III: Average horizontal errors using k-NN ( $k = 5$ ) with 150 m square size and different amounts of receivers. MLAT (5 sensors) is provided as comparison.

only “good” values below 10 into account, MLAT’s metrics are bound to improve vastly. However, doing this also decreases the number of usable messages by over 90%, reinforcing the fact that k-NN is useful in a much larger area.

The computational time is the trade-off for k-NN’s accuracy and robustness. Only with the largest square size of 600 m<sup>2</sup> it is in a range comparable to MLAT. However, depending on the density of the airspace and the available equipment, even larger grids and longer computation times would not pose a problem in real-world settings. The complexity of the algebraic MLAT algorithm is constant, while k-NN depends on the number of squares, i.e., both the size of the monitored area and the desired accuracy. If we assume a constant surveillance area, every reduction of the square size by half increases the number of squares, and thus the runtime, by four times. However, since a standard Macbook Pro (2013, 2.5Ghz i5, 8GB RAM) can run a single localization in less than 0.007 seconds on the 300 m<sup>2</sup> grid in our test area, the complexity remains negligible in ATC surveillance: more than 140 aircraft can be localized at the same time while maintaining an update rate of less than a second. In scenarios where localization is used mainly to verify suspicious aircraft data, the runtime of the algorithms is even less critical as the examined amount of data is small.

For our analysis, it is furthermore important to compare the impact of sensor numbers on localization. Table III shows the results for the same dataset and a 150 m<sup>2</sup> grid size, if only a subset of the five sensors receives the messages. After

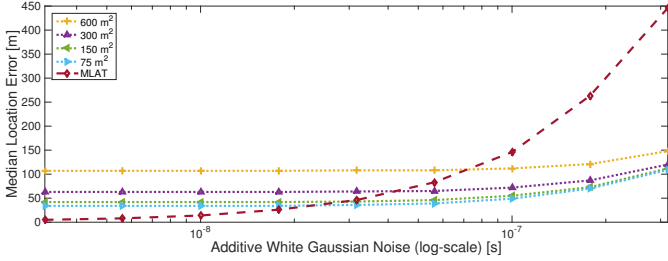


Figure 8: Median location error depending on the level of noise  $\sigma$  affecting the measurements. k-NN results obtained with  $k = 5$  and 4 different square sizes.

analysing all possible subsets and averaging the results, we conclude that with only three sensors sufficient horizontal accuracy can be achieved.

### B. Robustness

One of the main advantages of k-NN over MLAT is its robustness to noise. To analyze the impact of noisy signal measurements, we conduct simulations testing various noise levels against both algorithms, using a scenario similar to the one with which we collected our real-world data. We randomly distribute 5 sensors on a  $100 \times 100 \text{ km}$  grid, at altitudes randomly drawn from between 0 and  $1,000 \text{ m}$ . We create 100,000 aircraft signals sent from the grid at a height between  $10,000 \text{ m}$  and  $11,000 \text{ m}$ , calculate their ideal time of arrival at the 5 sensors and add white Gaussian noise  $Z_i \sim \mathcal{N}(0, \sigma)$  for each time  $i$ , where  $\sigma$  is the standard deviation given in seconds. We repeat this simulation 1,000 times at each  $\sigma$  level to smooth out effects of the sensor placement.

From the results shown in Fig. 8, we can conclude that the MLAT algorithm performs very well for no or very little noise. This is not surprising for an exact method. However, the algorithm is much less robust against increased noise levels compared to k-NN. As the noise approaches  $\sigma = 10^{-7.5} \text{ s}$ , we can see that the median location error of MLAT surpasses the one of k-NN with  $k = 5$  and  $75 \text{ m}^2 / 150 \text{ m}^2$  square sizes. This means that at a noise level of over  $31.6 \text{ ns}$ , our k-NN approach provides superior results.

Fig. 9 gives further insight into the underlying reasons for the performance difference between both algorithms. It shows the median location error depending on the GDOP of the measured signal. At a noise level of  $\sigma = 10^{-7} \text{ s}$ , we find that MLAT only performs well when GDOP is extremely low, while k-NN is unaffected by this problem. Comparing MLAT to k-NN with various grid densities, we can see that only signals with  $\text{GDOP} < 5$  can provide a low average localization error; higher GDOP levels are quickly outperformed by all square sizes. Unfortunately, such a low dilution of precision is present in only a very small fraction of the potential surveillance area (as illustrated in Fig. 4).

These findings illustrate that the MLAT algorithm requires extremely tight and costly synchronization and is still severely inhibited by the geometry of its receiver's locations. In contrast, k-NN with expected TDoAs can provide effective localization with good quality even with low-cost hardware




Figure 9: Median location error depending on the geometric dilution of precision affecting the measurements at a noise level of  $\sigma = 10^{-7} \text{ s}$ . k-NN results obtained with  $k = 5$  and 4 different square sizes.

such as Kinetic Avionics SBS-3 boxes and fully arbitrary receiver placement. These characteristics are very helpful in cases where deployment options are limited. Further, they can facilitate the use of TDoA-based localization in modern crowd-sourced receiver networks or pre-existing ADS-B deployments, both of which are regularly not optimized for perfect MLAT system geometry. It has to be noted that in professional ATC environments, continuous tracking algorithms further improve on the quality of the discrete localization results [21], however, this can be applied to all techniques equally.

## VIII. DATA VERIFICATION

We now consider the verification of ADS-B data using our approach. Data verification can be necessary for several reasons, such as transponder malfunctions, inaccurate aircraft instrument data, or the deliberate broadcast of false data. Without loss of generality, we assume such an adversarial case, where false location data is injected into the ATC system by using manipulated ADS-B messages. The attackers have two different mobility models, influencing the credibility of their positional claims. The data verification itself is done via a threshold: if an aircraft's ADS-B claim is deviating too much and for too long from the estimate, it is flagged as an anomaly, and thus as a possible attack.

### A. Test Data

We use our real-world flight data obtained from OpenSky to test our data verification scheme. Out of the whole dataset described in VI-A, we analyze 1,341 legitimate flights and show that they are accurately verified by our system. Furthermore, we use data from two simulated attacker types (due to ethical and legal reasons, we do not implement real-world attacks) on the ground and in the air and check whether they will be verified or not. Each attacker injects 200 messages with the legitimate coordinates of a real flight from our sample and follows specific location patterns:

- **Attacker 1** is ground-based and mobile, defined by a random horizontal start position on the grid, a random horizontal direction, and an altitude between 0 and  $500 \text{ m}$ , moving on the ground with a speed of  $50 \text{ km/h}$ .
- **Attacker 2** is a legitimate aircraft, broadcasting a wrong track. Its starting position is the same as the real aircraft



Attacker Type	Distance from claim [start/end/avg]
Ground, mobile	74.897 / 88.682 / 77.535 km
Aircraft	0 / 28.864 / 11.627 km

Table IV: Averaged horizontal distances from the two attackers’ real positions to their claimed aircraft positions during the time that flight data is injected.

but then real and fake track diverge horizontally at a random angle between 10 and 25 degrees (at cruising altitude), making attacker 2 more difficult to detect.

To illustrate their relative positions, Table IV shows the average deviations from the real track for both attacker types (at the beginning and the end of the attack, and on average, respectively). The attackers’ TDoAs are calculated by dividing the 3D distance between the sensors by the speed of light  $c$  and adding white Gaussian noise analogous to our real data to account for measurement and processing errors. We test each scenario 1000 times and analyze the detection rate.

### B. Evaluation

We now derive potential verification thresholds based on all legitimate flights and evaluate the resulting speed of detection against both attacker models. We further analyze how accurately the true origin of a signal can be located by our localization approach and MLAT.

Through experimental analysis of the legitimate data, we obtain a threshold that has not been exceeded by any of the legitimate flights in our dataset. More formally, we first obtain the number of consecutive messages  $n$  for given thresholds  $T$ , after which the number of failed verifications (or the probability of false alarms  $P_{fa}$ ) reaches 0. We then evaluate the time required to detect both attackers given  $T$  and  $n$ , such that the probability of missed detection  $P_{md}$  reaches 0, too.

Fig. 10 shows three thresholds  $T_1 = 250\text{ m}$ ,  $T_2 = 500\text{ m}$  and  $T_3 = 750\text{ m}$ , signifying the distance between our localization and the aircraft’s position claim, and their likelihoods  $P_{fa}$  of failing to verify a real flight depending on  $n$ . For threshold  $T_1$ , we find that  $P_{fa}$  reaches 0 when the average deviation between position claim and k-NN estimate exceeds 250 m over a contiguous period of 28 messages. For the higher thresholds  $T_2$  and  $T_3$ , respectively, 15 and 9 messages are sufficient.

Within these 28/15/9 messages, we detect all injections by attacker 1 (making  $P_{md} = 0$ ); their location always far exceeds all thresholds. Attacker 2, who starts out from the correct position of the impersonated aircraft, is much less obvious and continuously exceeds  $T_1$  after 41.7 messages on average (32.9 for  $T_2$ , 29.6 for  $T_3$ ). After 57 messages,  $P_{md}$  reaches 0, thus all injections have been detected (44 for  $T_2$ , 45 for  $T_3$ ).

This means, reliable detection is possible after about 22 seconds without any lost messages or 44 seconds assuming a typical 50% message loss on the ADS-B channel [22]. Naturally, the precise thresholds depend on equipment and scenario and should be trained and fitted accordingly.

Besides the simple verification of flight data, localization methods can provide an estimate of an attacker’s current location. Table V provides the results of this estimation for both attacker types. k-NN with expected TDoAs not only




Figure 10: Analysis of different thresholds. The graph shows the percentage of real flights that fail to verify as legitimate based on distance from the aircraft location claim and number of consecutive messages exceeding this distance. k-NN results obtained with  $k = 5$  and  $150\text{ m}^2$  square size.

Attacker Type	Distance to attacker [km]	
	k-NN	MLAT
Ground, mobile	1.918	44.947
Aircraft	0.145	0.270

Table V: Mean distances to actual horizontal location of an attacker. k-NN ( $k = 5$ ) with  $150\text{ m}$  square size.

accurately detects the distances between the attacker and the claim but can also give a good guess about the real origin of the signal. The horizontal estimate for the origin of message signals fits within approximately 2,000 m of the real location for the ground-based attacker type. For the aircraft attacker type, we obtain an estimate that is accurate within the typical range for legitimate flight localization as discussed in the previous section, i.e. with an error of less than 200 m.

In comparison, Table V illustrates the performance of MLAT in the same scenario. While MLAT can also verify the data and detect deviations between the signal origin and the position claim, it can not provide an accurate guess of the ground-based attackers type. Caused by large GDOP values of 1000 and more, the sensors’ estimates of other objects in the same plane (i.e., on or near the ground) suffer from an average error of more than 44 km, making them essentially unusable. The origin of the aircraft attacker 2 is estimated with a low error of 270 m, which is in line with MLAT’s performance for the localization of legitimate flight data in our setup.

This shows that, while it is entirely feasible (though costly) to build an MLAT system with good accuracy even for large surveillance areas in high altitudes, it is difficult to provide the same level of accuracy on the ground with the same system, especially when using cheap and unplanned deployments.

## IX. CONCLUSION

We proposed a new method for the localization of aircraft based on the expected time differences of arrival and the k-Nearest Neighbor algorithm. We evaluated our scheme with real-world flight data from our large-scale sensor network OpenSky using only low-cost ADS-B sensors in an unplanned and crowdsourced deployment. We find that it outperforms the popular multilateration approach in terms of range and coverage. Since it is not similarly affected by dilution of

precision, it is possible to increase the overall coverage by a factor of up to 2.5 in our examined setting.

In terms of accuracy, our results show that the mean aircraft location accuracy can be increased by up to 41% in comparison with an algebraic MLAT algorithm. Furthermore, as it does not suffer from dilution of precision to the same extent, our approach is also more robust and less susceptible to noisy environments and bad system geometry, making it a better choice for pre-existing and unplanned receiver deployments.

Finally, we compare our approach against MLAT for the verification of aircraft location claims. The increased coverage improves the verification radius and enables us to detect even sophisticated aircraft-based attackers within 40 seconds. Furthermore, contrary to MLAT, it is possible to detect the approximate location of ground-based attackers within a mean horizontal error of about 2,000 m.

Further work will involve analyzing the capabilities of our approach when it comes to localization in 3D space. While 3D localization is a natural extension of this work, it is significantly more difficult as the degrees of freedom are increased. Nevertheless, even a less accurate solution could be helpful for verification purposes, in particular compared to MLAT, which is generally not able to give any useful altitude approximations of aircraft [23].

## REFERENCES

- [1] Boeing, “Boeing current market outlook 2015-2034,” Tech. Rep., 2015. [Online]. Available: <http://www.boeing.com/commercial/market/>
- [2] US Department of Transportation, “Unmanned aircraft system (uas) service demand 2015- 2035,” Tech. Rep., 2013. [Online]. Available: <https://fas.org/irp/program/collect/service.pdf>
- [3] M. Strohmeier, M. Schäfer, R. Pinheiro, V. Lenders, and I. Martinovic, “On perception and reality in wireless air traffic communication security,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 6, pp. 1338–1357, Jun. 2017.
- [4] V. Bryan and P. Maushagen, “Flightradar24 finds not just planespotters flocking to its website,” *Reuters*, May 2015.
- [5] M. Strohmeier, M. Smith, M. Schäfer, V. Lenders, and I. Martinovic, “Assessing the Impact of Aviation Security on Cyber Power,” in *8th International Conference on Cyber Conflict (CyCon)*. NATO CCD COE, 2016, pp. 223–241.
- [6] M. Strohmeier, V. Lenders, and I. Martinovic, “Lightweight location verification in air traffic surveillance networks,” in *Proceedings of the 1st ACM Workshop on Cyber-Physical System Security (CPSS)*. ACM, 2015, pp. 49–60.
- [7] E. A. Lester, “Benefits and incentives for ads-b equipage in the national airspace system,” Ph.D. dissertation, Massachusetts Institute of Technology, 2007.
- [8] M. Schäfer, M. Strohmeier, V. Lenders, I. Martinovic, and M. Wilhelm, “Bringing Up OpenSky: A Large-scale ADS-B Sensor Network for Research,” in *Proceedings of The 13th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, Apr. 2014, pp. 83–94.
- [9] H. Liu, H. Darabi, P. Banerjee, and J. Liu, “Survey of wireless indoor positioning techniques and systems,” *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, vol. 37, no. 6, pp. 1067–1080, 2007.
- [10] P. Bahl and V. N. Padmanabhan, “RADAR: An in-building RF-based user location and tracking system,” in *Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, vol. 2. IEEE, 2000, pp. 775–784.
- [11] B. Li, Y. Wang, H. K. Lee, A. Dempster, and C. Rizos, “Method for yielding a database of location fingerprints in wlan,” *IEEE Proceedings-Communications*, vol. 152, no. 5, pp. 580–586, 2005.
- [12] C. Ning, R. Li, and K. Li, “Outdoor location estimation using received signal strength-based fingerprinting,” *Wireless Personal Communications*, pp. 1–20, 2016.
- [13] A. Rozyyev, H. Hasbullah, and F. Subhan, “Combined K-Nearest Neighbors and Fuzzy Logic Indoor Localization Technique for Wireless Sensor Network,” *Research Journal of Information Technology*, vol. 4, no. 4, pp. 155–165, 2012.
- [14] I. A. Mantilla-Gaviria, M. Leonardi, G. Galati, and J. V. Balbastre-Tejedor, “Localization algorithms for multilateration (mlat) systems in airport surface surveillance,” *Signal, Image and Video Processing*, vol. 9, no. 7, pp. 1549–1558, 2015.
- [15] M. Mosavi and H. Azami, “Applying Neural Network Ensembles for Clustering of GPS Satellites,” *International Journal of Geoinformatics*, vol. 7, no. 3, pp. 7–14, 2011.
- [16] R. B. Langley, “Dilution of precision,” *GPS world*, vol. 10, no. 5, pp. 52–59, 1999.
- [17] N. J. Gomes, P. P. Monteiro, and A. Gameiro, *Next generation wireless communications using radio over fiber*. Chichester, UK: John Wiley & Sons, 2012.
- [18] M. Schäfer, M. Strohmeier, M. Smith, M. Fuchs, R. Pinheiro, V. Lenders, and I. Martinovic, “Opensky report 2016: Facts, figures and trends in wireless atc communication systems,” in *35th IEEE/AIAA Digital Avionics Systems Conference (DASC)*, Sep 2016.
- [19] M. Strohmeier, I. Martinovic, M. Fuchs, M. Schäfer, and V. Lenders, “OpenSky: A Swiss army knife for air traffic security research,” in *34th IEEE/AIAA Digital Avionics Systems Conference (DASC)*. IEEE/AIAA, Sep 2015, pp. 1–14.
- [20] A. C. Andersen, “Comparative analysis of multilateration methods for signal emitter positioning,” Tech. Rep., Jun. 2011.
- [21] S. S. Blackman, *Multiple-target tracking with radar applications*. Artech House, Inc., Dedham, MA, 1986, vol. 1.
- [22] M. Strohmeier, M. Schäfer, V. Lenders, and I. Martinovic, “Realities and Challenges of NextGen Air Traffic Management: The Case of ADS-B,” *IEEE Communications Magazine*, vol. 52, no. 5, pp. 111–118, 2014.
- [23] G. Galati, M. Leonardi, P. Magarò, and V. Paciucci, “Wide area surveillance using SSR mode S multilateration: advantages and limitations,” in *European Radar Conference (EURAD)*, 2005, pp. 225–229.



**Martin Strohmeier** (martin.strohmeier@cs.ox.ac.uk) is a postdoctoral researcher in the Department of Computer Science, University of Oxford. His current research interests are mostly in the area of wireless security and air traffic communication. Before coming to Oxford, he received his MSc degree from TU Kaiserslautern, Germany and worked as a researcher at Lancaster University’s InfoLab21 and Deutsche Lufthansa AG.



**Vincent Lenders** (vincent.lenders@armasuisse.ch) is a research director at armasuisse. He leads the cyber and information sciences research activities of the Swiss Federal Department of Defense. He received the MSc and PhD degrees in electrical engineering from ETH Zurich. He was a postdoctoral research fellow at Princeton University. He is a co-founder and board member of the OpenSky Network and Electrosense Associations. His current research interests are in the fields of cyber security, information management, big data, and crowdsourcing.



**Ivan Martinovic** (ivan.martinovic@cs.ox.ac.uk) is an Associate Professor at the Department of Computer Science, University of Oxford. Before coming to Oxford, he was a postdoctoral researcher at the Security Research Lab, UC Berkeley and the Secure Computing and Networking Centre, UC Irvine. He obtained his PhD from TU Kaiserslautern and MSc from TU Darmstadt, Germany.